Drivers-Inspired Ants for Solving the Vehicle Routing Problem with Time Windows

Dalicia Bouallouche* and Jean-Baptiste Vioix†
LE2I, UMR6306 (CNRS),
University of Bourgogne Franche-Comté.
Auxerre France.
Email*: dalicia.bouallouche@u-bourgogne.fr
Email†: jean-baptiste.vioix@u-bourgogne.fr
Tel.*: +33-03-86492846
Tel.†: +33-03-86492841

Eric Busvelle
Aix Marseille Univ,
Université de Toulon,
CNRS, ENSAM, LSIS,
Marseille, France.
Email:busvelle@univ-tln.fr
Tel.: +33-04-94146766

Stéphane Millot
Tedies,
40 Rue Louis Richard
89000 AUXERRE France.
Email:stephane.millot@tedies.fr
Tel.: +33-03-86539834

***Abstract*—In our study, we develop a method that merges two information sources within ants colony optimization heuristic. Namely artificial ants which occurs for short term optimization and transporter's vehicles that occurs in long term and continuous optimization toward solving the real-world vehicle routing problem. This study is supported by a transporter (Upsilon) of the region of l'Yonne in France and a transport and logistics software development company (Tedies). Our method suits for transporters that use human planners to make decisions about their tours and intending to move to computer planners without drastically upsetting the drivers habits. Hence, the pledge of this study is to take advantage from transport operators practices to achieve solutions which are as close as possible to the real-world vehicle routing planning, and keep a human control on the way optimal paths are computed and applied.**

I. Introduction

In this paper, we propose a new approach for solving the VRP (Vehicle Routing Problem). Our aim is to inspire the ACO (Ant Colony Optimization) with the real-world transporter practices that include the transporter employees field experience, particularly, the planners and the drivers. Our study builds upon real-world data of a carrier, Upsilon, and a company, Tedies that proposes transport and logistics software for carriers, both operating in the region of lYonne, France. Doing so, we benefit from the experience that people acquire during their professional career, including the insight and intelligence of human resources.

We solve our vehicle routing problem with a cluster-first route-second approach [1]. To be accomplished, both of these two stages take advantage from the transporters real-world practices with the objective of finding more familiar solutions to the drivers habits. Another purpose of our work is to update information of the road network database, especially concerning trucks routes specificities. This is accomplished by retrieving the feedback after practical routes made by the fleet of the vehicles.

The solution that we develop is supposed to accompany the transport operators and the drivers. Thus, we propose to the transporter a method that gives a solution for its vehicle routing problem with the advantage that the resulting routes could be modified by the transport planners and the drivers if it provides an improved solution. The solution routes including the modifications when put into practice will be used as a feedback for next optimization or simply update and complete the map database. It is a way to earn the trust of the carrier by offering a customizable system and by taking into consideration the decisions of more experienced people within our system. In addition, it allows the product to be more easily accepted by end users who are drivers.

Our proposed approach deals with an offline optimization problem. Since our system updates the data from new trips borrowed by the drivers after they back to the depot, then the updates will be considered for next optimizations. Also, in our case, the information data exchange (related to customers demands) arrives early before the trucks start their routes (before 3:00am at Upsilon), we run the routes planning with the data received at this moment, as well as the data of permanent customers demands, and no new demand is taken into account during the vehicle routes, then, an online routing update is not needed in this case.

In the next session, we will give a short review of the vehicle routing problem. In section III, details about our approach are stated. Then, we will move to the presentation of some actual results in section IV. Finally, we will give a conclusion and short-term perspectives in section V.

II. Background and discussion

The vehicle routing problem can be defined as : *given a set of customers requesting for a delivery service to a transporter, these customers have to be visited once (and only once) by one vehicle of the transporter fleet and the routes must start at the depot. This must be achieved by minimizing a cost which can be, for instance, the total amount of time spent on the road by each driver.* It is an NP-hard combinatorial optimization problem.

Since the sixties, several research studies have been concentrated on solving the VRP and its variants with various methods. Meta-heuristic methods can give good quality solutions in a fairly short time [2]–[4]. The most common VRP variant is the VRPTW (Vehicle Routing Problem with Time Windows) which is the VRP where the vehicles have to deliver each customer within its specified time window and the total weight of each vehicle including all the customers demands (parcels) that it will deliver must not exceed its capacity. We refer the reader to [5] and [6] for the VRPTW problem formulation and a survey of the VRPTW resolution methods. Other VRP variants formulations, overview of exact and heuristic VRP resolution methods and it variants as well as real-world VRP applications, technologies and software can be found in [7].

Theoretical VRP problems, for which authors propose methods for solving VRP problems with fixed parameters [6] are not efficient for real-world instances. Then, researchers interested on the real-world VRP category, and develop more adapted methods for real-world situations [8]. This VRP variant is referred to as RVRP (Rich Vehicle Routing Problems) [9]. Generally in this case, some necessary parameters to the resolution of the VRP could be unknown or inaccurate. For instance, we can mention travel times, service times, errors in customers address geocoding [10], new or evolving customers demands during the problem resolution. In [8], authors propose a survey of hybrid solution methods for the RVRP. Various studies have been devoted to solve real-life variant of the VRP. In the following, we will cite some research studies that use ant colony optimization system to deal with the real-world VRP. In [3], the authors work on two real-world industrial-scale applications for solving two real-world VRP variants. They adapt the MACS-VRPTW algorithm [11] for real-world applications for a major supermarket chain in Switzerland to solve a VRP with time windows problem instance and for a leading distribution company in Italy to solve a VRP with pickup and delivery problem. In [12], Montemanni et al. apply an ACS (Ant Colony System) based approach on a realistic case study which deals with fuel distribution company data in the city of Lugano in Switzerland to validate the convenience of the method on real-life situations.

The goal of our study is to provide a practical tool which relies on a mutual benefit between the decisions made by our algorithm and those of the transporter to solve the VRPTW. This tool is created for Tedies to provide routing planner software that suits the expectations of the carriers, that is to say, a more realistic tool.

We have noticed, through our conversations with Upsilon, that deploying an algorithm that suits the drivers habits is a crucial point. The same ascertainment have been mentioned in the study of Rizzoli et al. [3]. Indeed, we must not drastically shake up the drivers habits. Rather, our mission is to integrate in their mind right reflexes through the feedback. Their practices will then influence our algorithm toward good quality solutions. From this, we focused on developing a tool that proposes permanent and automatic interaction between professional experience and optimization algorithm, ensuring the reliability of optimal solutions. For this study, we do not take into consideration all the transporters constraints, but we restrict our study on solving the VRPTW for keeping only the main characteristics which are useful for carrying out our objective.

Through our literature reviews, we have noticed that the ACO has proved its efficiency to real-life vehicle routing problems [3]. Also, the ACO fits better with the approach that we seek to implement. In fact, the data of real vehicles itineraries, which is our feedback data, is injected in the ants pheromone trail of the ACO to influence the ants by the human planners and the drivers decisions. This justifies the choice of using ACO for solving our problem.

We will focus on solving the TSPTW (Traveling Salesman Problem with Time Windows) with an ant colony optimization based resolution method because, in this paper, we deal with a cluster-first route-second strategy. To solve the TSP with ACO, a set of artificial ants have as task to find the best (shortest or fastest) vehicle route. Each one passes through the graph of customers to search for a path (as stated in the first paragraph of this section), starting from the depot, then stochastically and progressively choose customers to add to the partial path and finally back to the depot [13]. They deposit pheromone trail which dynamically evolves during solution construction according to the quality of the solution (path) found by each ant, which are in turn influenced by the pheromone trail.

We refer the reader to the section 2 of the paper of Stützle et al. [14] for more details on solving the TSP with ACO heuristics.

III. General approach

The vehicles trips of the transporter are tracked through Tomtom® GPS units and the corresponding data (recording time, GPS coordinates, speed, ...) is recorded with a frequency of 0.1 Hz. We can deduce travel times of the road segments (ways) from the vehicles speeds saved during the passages of the vehicles through the ways of the road network. If a way has many records, we take the average speed of this way. The graph of customers is deduced from the road network. In fact, having a set of customers to deliver, we construct a complete and asymetric graph, where the costs of the edges are the travel times computed from the fastest paths going from each customer to another. We use the road network of OpenStreetMap®, on which we select the region of Burgundy (France) where Upsilon perform its routing process.

Figure 1 shows the general steps for solving our vehicle routing problem by integrating the real routes feedback. We first notice that the Upsilon's fleet of vehicles serves almost the same region and the majority of the transporter customers are permanent. Only 30% of new customers apply for collection/delivery requests. The habitual and daily process of the transport operators is to dispatch the customers requests among the vehicles, then the experienced drivers planify the visiting order of customers. The real vehicle paths that have been recorded by Tomtom GPS units during their routes will be used to help the ACO heuristic and inspire it by a good and realistic practical solution to enhance the solution quality of our next optimization problems and make the solution converge faster.

For our clustering, we develop a carrier real-world inspired clustering which is motivated by geographical and practical considerations. The transport planners naturally follow division by town areas, which is an intuitive way to construct clusters as the customers are mainly grouped into clusters in different towns. Our aim is to draw upon the practices of experienced transport professionals to construct our clusters. We will not detail our clustering algorithm in this paper because we want to emphasize our contribution to the ant colony optimization heuristic. However, to provide insights, we dispatch the customers requests into clusters according to the historical dispatch of the vehicle routes of the carrier, in order to get clusters that matches the geographical division of regions. Other clustering methods can be adopted, however, to go with the rest of our heuristic, the clustering method have to give results that suits the drivers habits and each cluster (vehicle) have to evolve daily in the same area.

The clustering is followed by finding an optimal route on each cluster of customers. It consists of the TSPTW, that we solve with an ant colony optimization metaheuristic that we have adapted to deal with our objective.

The drivers are not supposed to apply our solution exactly as it is provided, they are free to make some alterations if they guess that it may enhance the solution quality in terms of cost and/or more realistic tours completion that suits the ground reality (taking a fastest path than the one given by our system, bypass work areas recently identified by the drivers, exchange two delivery points between two routes, ...). This strategy will enhance the database with new information without corrupting the optimal solution if the driver is wrong. The drivers will then go for deliveries following our optimized routes which include their personal touch. Then the real truck routes will be saved through the GPS units so as to be used as a feedback for future optimizations. We integrate the trucks feedback data in the ACO heuristics, through the pheromone trail of the ants that we combine with the trucks trail. This includes a step of segmenting the real routes to reconstitute the route edges and then add trucks trail to the initial amount of pheromone, corresponding to the real routes edges. During the solution construction, we also consider the trucks trail in the step of pheromone trail update, in order to keep the influence of real-life vehicle routes. The details on the real routes segmentation, pheromone initialization and pheromone update are stated in sections III-A, III-B and III-C, respectively.

In addition to travel times update, the feedback data can give us information on prohibited access to trucks according to the vehicle category, also, we can deduce the service times at customers by identifying the breakpoints duration at each customer during the route.

To summarize, it is an optimization loop that contains three main steps: solution construction (with ACO and transport operators experience), practical implementation of the solution and using the feedback for enhancing the realizability and the quality of the next problem solution.

A. Segmentation of real vehicle routes

This step consists of real (practical) route segmentation in order to reconstitute the edges that compose each real

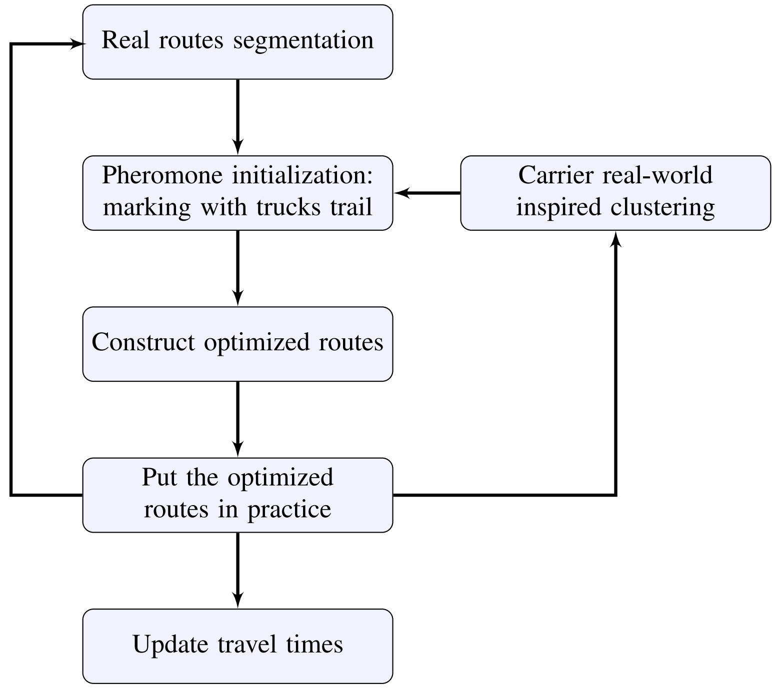

Fig. 1. Trucks routes feedback integration to solution.

path borrowed by a vehicle. The extremities of the edges are determined by the visited customers. The customers locations are known by our system, but since the visiting order of customers in practice can be modified by the drivers we have to re-identify them on the vehicle itinerary (because the practical route would be different from our computed routes), . On the real path of the vehicle, the visited customers are known as collection/delivery points. We localize the delivery points on a vehicle path where the speed of the vehicle is null, then we associate each delivery point to a customer. We refer the reader to [10] for more details about delivery points identification. Figure 2 shows an example of the segmentation on a vehicle path. The points are the customers (delivery points) and the edges resulting from segmentation are delimited between each pair of successive customers in the vehicle route.

B. Pheromone initialization using trucks trail

By this step we want to incorporate the trucks trail in the ants pheromone trail of the ACO heuristics. The trucks trail is used to inspire the ants by real-life practices and the pheromone trail is used to guide the ants to move toward an optimal solution. The ACO heuristics in our study will contain information about the historical real vehicles paths by adding an amount of trucks trail, referred to as ”vehicles (trucks) pheromone trail”, according to the costs (real travel times) of the edges resulting from routes segmentation. In other terms, we add to the initial ant pheromone trail, the vehicles pheromone trail on the graph edges that have been borrowed by these vehicles during the past real routes. The quantity of the vehicles pheromone trail that we add to the

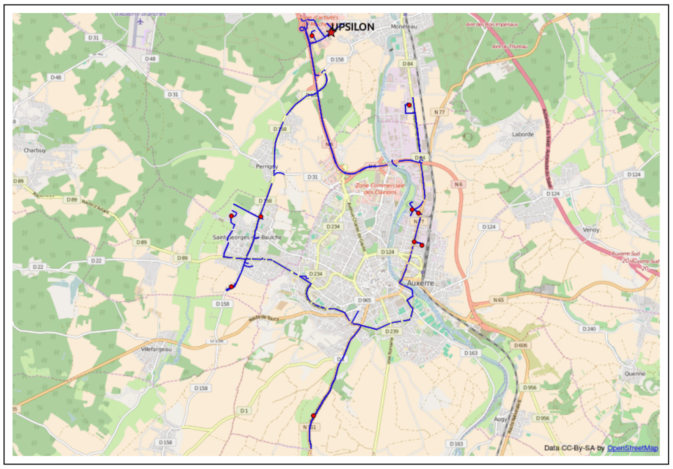

Fig. 2. Edges reconstitution from a vehicle path on customers locations

corresponding edges depends on the real routes costs. We, hence, proceed to the pheromone initialization as follows.

$$\tau_0(e) = \begin{cases} \tau_0 & \text{if } e \in E \setminus U \\ \tau_0 + \lambda \times \left(\frac{1}{\mathcal{C}_R} + \frac{1}{\mathcal{C}(e)}\right) & \text{if } e \in U \end{cases}$$

Such that U is the set of graph edges resulting from the segmentation of the vehicle paths (in section III-A). τ_0 is the default initial amount of pheromone. $\mathcal{C}_R$ is the total travel time of the real vehicle route R and $\mathcal{C}(e)$ is the travel time of the edge e that belongs to R. The rate λ defines the intensity of the trucks trail. $\tau_0(e)$ is the initial amount of pheromone of edge e including, eventually the trucks trail.

The edges that have at least an extremity which is a new customer will only be processed with default initial pheromone initialization (with value τ_0) since the customer does not exist in previous routes. This does not impact the solution because, during the solution computation, the ants of the ACO heuristic will add these new customers to the vehicles routes (under construction) at the best positions.

C. Pheromone update

In the pheromone update step of the ACO heuristic, we will update the quantity of pheromone according to the paths found by the elite ants.

First, the pheromone evaporation is processed on all the edges of the graph G, as follows.

$$\tau(e) = (1 - \rho)\tau(e) + \rho\tau_0(e) \tag{1}$$

where ρ is the rate of pheromone evaporation and, instead of using default pheromone value (in the second term of the equation (1)) as for a basic pheromone evaporation, we use $\tau_0(e)$, which is the initial amount of pheromone on edge e which, eventually, includes the vehicles pheromone trail computed in sections III-A and III-B.

Second, through the following equation we show how the amount of pheromone is updated by adding pheromone trail to the graph according to the paths costs of the elite ants.

$$\tau(e) = min(\tau(e) + \frac{\sigma - \mu}{\mathcal{C}_\mu}, \tau_{max})$$

TABLE I
SOME VRP RESULTS

Day	HumPlan(h)	HeurOpt(h)	Gain(h)	NbCust	Runtime(m)
1	24.81	13.8	11.01	98	4.37
2	20.38	12.87	7.51	85	3.43
3	22.71	14.55	8.16	96	3.93
4	19.34	13.44	5.90	70	2.31
5	22.91	16.74	6.17	103	4.39
6	18.68	10.39	8.29	71	2.99
7	20.78	12.46	8.32	72	2.80
8	20.20	12.9	7.30	88	3.39
9	18.37	11.41	6.96	79	3.35
10	18.25	14.98	3.27	109	5.13
11	23.45	19.43	4.02	117	6.55
12	17.75	14.11	3.64	88	4.22
13	24.89	15.42	9.47	92	3.84
14	20.05	16.02	4.03	104	5.46
15	26.94	20.99	5.95	128	6.00
AVG	21.30	14.63	6.67	94	4.14

where σ is the number of elite ants, μ is the rank of the elite ant (best ant is of rank 0) and $\mathcal{C}_\mu$ is the total time cost of the elite ant μ path. The maximum value of the pheromone can not exceed τ_{max} to allow more diversification by avoiding the creation of highways of pheromone to which the ants will be easily attracted [15]–[17]. Then, $\tau(e)$ is the amount of pheromone on edge e at the current ACO iteration.

IV. EXPERIMENTATION, RESULTS AND DISCUSSION

We work on a fleet of eight vehicles through which we have collected around seven months of data. We compare the solution resulting from our algorithm to the routes planned by human planners of Upsilon.

In our algorithm, the service times at the customers are not considered, as we set them to null value. This is why for our comparisons with real vehicle paths set in tour, we have deleted from the vehicles total traveling times the periods during which the vehicles were at standstill to get only the time during which the vehicles were moving. Although we have the real travel durations of the vehicles, we judged that it is more convenient to use our processed data (travel times) retrieved from all the vehicles through all the days of recording (see first paragraph of section III). The results of table I are obtained by setting $\tau_0 = 1$, $\tau\_max = 20$ and $\rho = 0.7$, for 100 iterations and a number of ants and elite ants respectively set to half and to quarter of the number of customers. Details about parameter setting of ACO heuristics can be found in [14]. The value of parameter ρ might be rather high in order to ensure the persistence of the initial pheromone containing the trucks trail. This occurs in equation (1) where 70% of the amount of pheromone will be evaporated and replaced by 70% of the initial amount of pheromone containing the trucks trail. It leads to more diversification as well as ensuring the trucks trail persistence through the iterations.

Note that the paths on which we conduct the vehicles could be forbidden for lorries. In fact, we have no information about ways with denied access to heavy trucks in OpenStreetMap® database. However, we can retrieve the feedback from the real vehicle routes accomplished by the drivers to progressively complete OpenStreetMap® database with denied access according to the vehicles kind.

Including the trucks trail in the pheromone initialization is a way to start from feasible paths and give to the heuristics a high probability to improve the quality of the real human planned routes, and reintroducing the trucks trail in the pheromone update with a given rate drives the heuristic toward the right direction, and avoid constructing more expensive routes than human planners (It is what we will see in figure 3). If not, all the initial amount of pheromone which contains trucks trail will disappear through the iterations because of the evaporation step of the pheromone update.

We have compared in table I human planned vehicle routes (HumPlan) to our computed vehicle routes (HeurOpt), with the corresponding runtime which does not include the clustering computation. The routes costs in the table are in hours and the runtime in minutes. Gain is the vehicle routes traveling time that we have saved by using our heuristic. The number of customers delivered each day by the fleet of the vehicles is indicated in column $NbCust$. We have an average gain of 6.67 hours per day for the whole fleet in 4.14 minutes of runtime.

Figure 3 shows the convergence of the solutions costs in function of the number of iteration in the ACO heuristic, in the situation where we have injected the trucks trail in the ants pheromone trail and in the situation where not (by default only ants pheromone). For this test, we deal with 15 days of tours of a vehicle, the solution cost in the figure is the mean of the costs of all these tours, given an iteration. We clearly see a gap between the two curves as the solution costs with trucks pheromone trail has a head start given by the historical routes of the vehicle, thanks to the pheromone initialization containing trucks trail, and this gap is kept throughout the ACO iterations, thanks to our pheromone update step that maintains the influence of the trucks trail.

Under this feedback loop system, namely the continuous re-injection of the optimization results of the last vehicle routes for solving the next vehicle routing problem of the transporter, the ACO heuristic will always have a good depart, as the ants will be inspired by the trucks trail all along searching for a solution. Thus, the solution will be improved in permanence through the days.

V. Conclusion and perspectives

This paper focuses on how we make human planners and drivers experience cooperate with the ants of the ACO heuristics in order to solve the real-world VRPTW.

We use ants for their ability to explore systemically a lot of solutions but we use also the experience of drivers to help the algorithm to reach more quickly a high quality solution and to enrich the database with new paths. Ants and human are highly complementary.

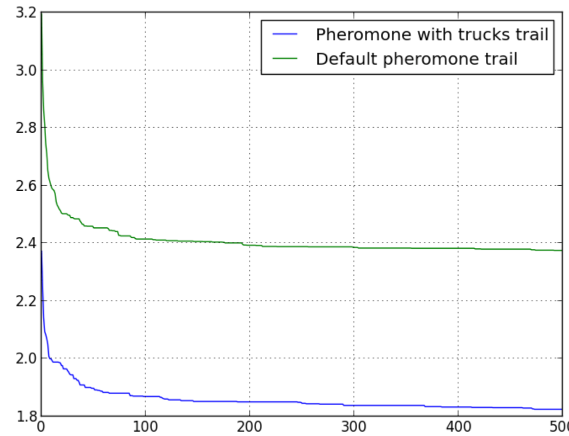

Fig. 3. Routes cost convergence through the ACO iterations

Our aim with this approach is not to instantly reach the best solution as soon as we start optimizing the vehicle routes of the transporter. But, we propose an optimization loop that includes what we have learned from the last practical routes in the future vehicle routing optimizations, then, the quality of the solutions will be progressively enhanced through the days. We deal with a classical vehicle routing problem with time windows and backhauls by leaving aside other real-world parameters to concentrate on exposing the idea of introducing the real-world routes feedback in the problem resolution with ant colony optimization. We note that our method is suitable for transporters which have regular customers.

In future works, we plan to develop more deeply our methods, concerning mainly the clustering method, the details and other characteristics of our ACO heuristics and emphasize other findings. In addition, we will integrate this tool in Tedies software and run it with Upsilon data in order to see the evolution of the quality of the vehicle routes through the days.

Acknowledgment

We thank the personnel of Upsilon and Tedies for their help in order to carry our experimentations.

References

[1] D. T. HIQUEBRAN, A. S. ALFA, J. A. SHAPIRO, and D. GITTOES, "A REVISED SIMULATED ANNEALING AND CLUSTER-FIRST ROUTE-SECOND ALGORITHM APPLIED TO THE VEHICLE ROUTING PROBLEM," *Engineering Optimization*, vol. 22, no. 2, pp. 77–107, dec 1993. [Online]. Available: http://dx.doi.org/10.1080/03052159308941327

[2] C. Blum and A. Roli, "Metaheuristics in combinatorial optimization: Overview and conceptual comparison," ACM Comput. Surv., vol. 35, no. 3, pp. 268–308, Sep. 2003. [Online]. Available: http://doi.acm.org/10.1145/937503.937505

[3] A. E. Rizzoli, R. Montemanni, E. Lucibello, and L. M. Gambardella, "Ant colony optimization for real-world vehicle routing problems," Swarm Intelligence, vol. 1, no. 2, pp. 135–151, 2007.

[4] R. F. Hartl, G. Hasle, and G. K. Janssens, "Special issue on rich vehicle routing problems," Central European Journal of Operations Research, vol. 14, no. 2, pp. 103–104, 2006.

[5] J.-F. Cordeau, G. Desaulniers, J. Desrosiers, M. M. Solomon, and F. Soumis, “7. VRP with time windows,” in *The Vehicle Routing Problem*. Society for Industrial & Applied Mathematics (SIAM), jan 2002, pp. 157–193. [Online]. Available: http://dx.doi.org/10.1137/1.9780898718515.ch7

[6] O. Bräysy and M. Gendreau, “Vehicle routing problem with time windows, part i: Route construction and local search algorithms,” Transportation science, vol. 39, no. 1, pp. 104–118, 2005.

[7] P. Toth and D. Vigo, *Vehicle routing: problems, methods, and applications*. Siam, 2014, vol. 18.

[8] K. F. Doerner and V. Schmid, *Survey: matheuristics for rich vehicle routing problems*. Springer, 2010.

[9] R. Lahyani, M. Khemakhem, and F. Semet, “Rich vehicle routing problems: From a taxonomy to a definition,” European Journal of Operational Research, vol. 241, no. 1, pp. 1–14, 2015.

[10] D. Bouallouche, J.-B. Vioix, S. Millot, and E. Busvelle, “*Address Matching Using Truck Tours Feedback*,” in Operational Excellence in Logistics and Supply Chains. Proceedings of the Hamburg International Conference of Logistics (HICL)22. *Hamburg, Germany*, T. Blecker, W. Kersten, and C. M. Ringle, Eds., vol. 22. epubli, Sept 2015, pp. 391–418.

[11] L. M. Gambardella, É. Taillard, and G. Agazzi, “Macs-vrptw: A multiple ant colony system for vehicle routing problems with time windows,” Tech. Rep., 1999.

[12] R. Montemanni, L. M. Gambardella, A. E. Rizzoli, and A. V. Donati, “Ant colony system for a dynamic vehicle routing problem,” Journal of Combinatorial Optimization, vol. 10, no. 4, pp. 327–343, 2005.

[13] M. Dorigo and T. Stützle, “Ant colony optimization: overview and recent advances,” in Handbook of metaheuristics. Springer, 2010, pp. 227–263.

[14] T. Stützle, M. López-Ibánez, P. Pellegrini, M. Maur, M. M. De Oca, M. Birattari, and M. Dorigo, “Parameter adaptation in ant colony optimization,” in Autonomous search. Springer, 2011, pp. 191–215.

[15] C. Solnon, *Optimisation par colonies de fourmis*. Hermès science publ.-Lavoisier, 2008.

[16] W.-C. Xie, C.-M. Chen, S.-S. Fan, and L.-L. Li, “A vehicle routing optimization method with constraints condition based on max-min ant colony algorithm,” Appl. Math, vol. 8, no. 1L, pp. 239–243, 2014.

[17] M. Dorigo, M. Birattari, and T. Stutzle, “Ant colony optimization: artificial ant as a computational intelligence technique. university libre de bruxelles,” IRIDIA Technical report Series, Belgium, Tech. Rep., 2006.